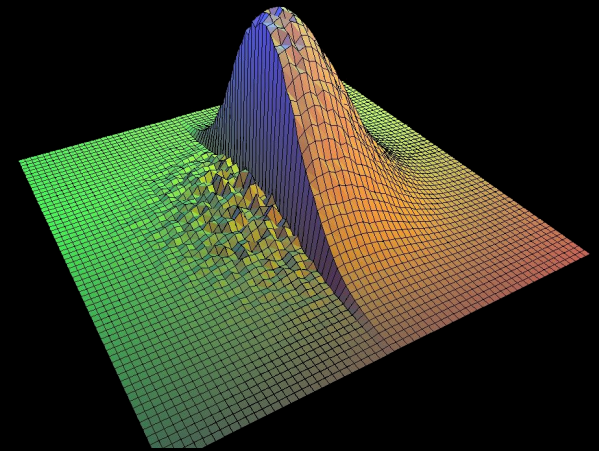


Fast Bilateral Filtering for the Display of High-Dynamic-Range Images

Frédo Durand & Julie Dorsey
Laboratory for Computer Science
Massachusetts Institute of Technology

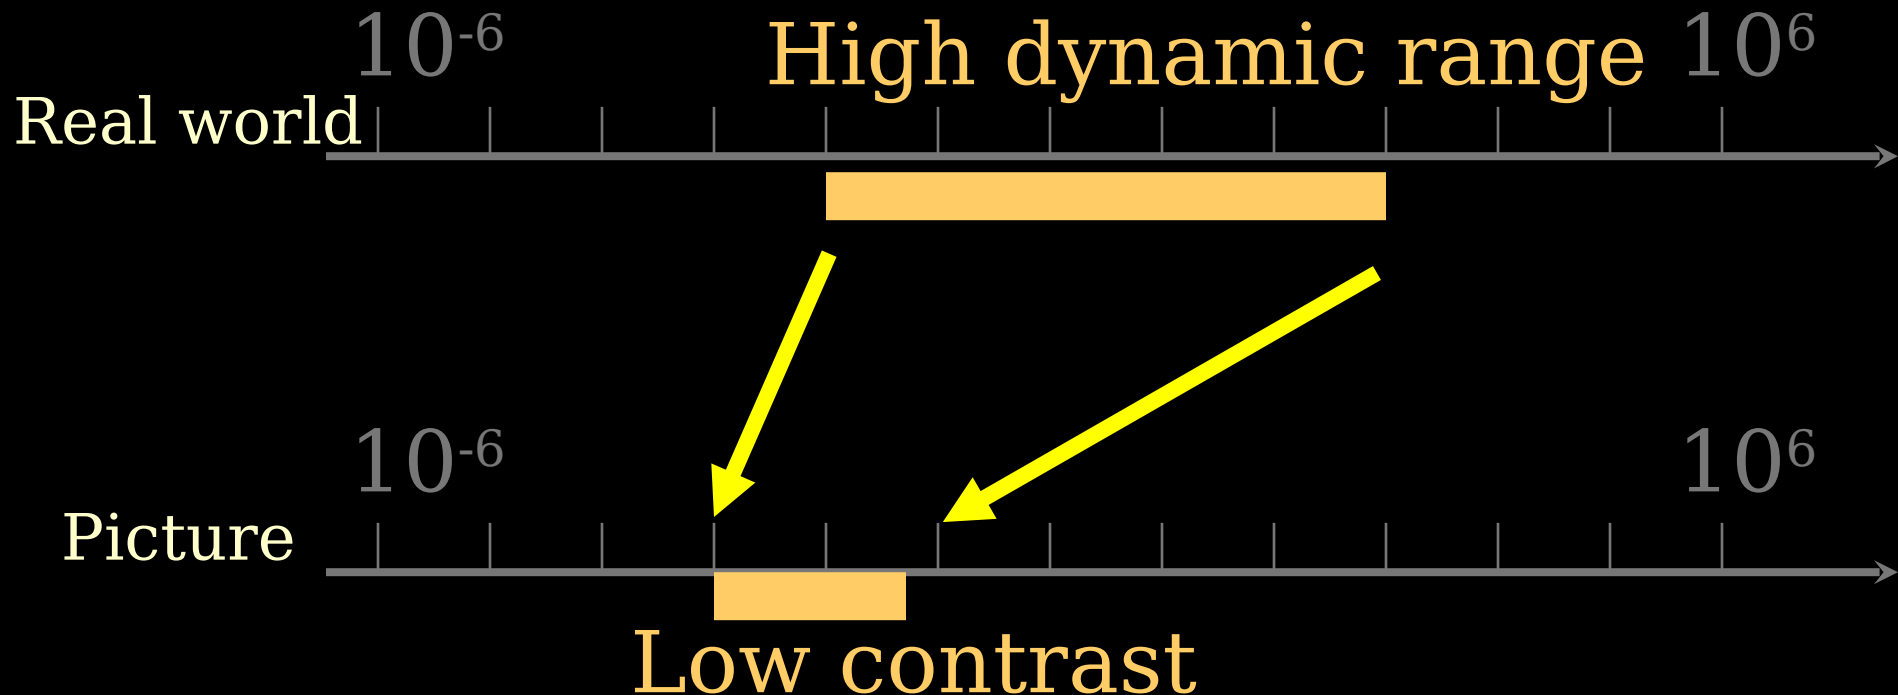
Contributions

- Contrast reduction for HDR images
- Edge-preserving filter



Contrast reduction

- Match limited contrast of the medium
- Preserve details



A typical photo

- Sun is overexposed
- Foreground is underexposed



Gamma compression

- $X \rightarrow X^\gamma$
- Colors are washed-out

Input



Gamma



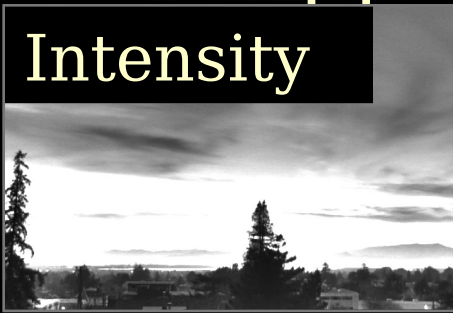
Gamma compression on

~~intensity~~

Colors are OK,

but details (intensity high-frequency)

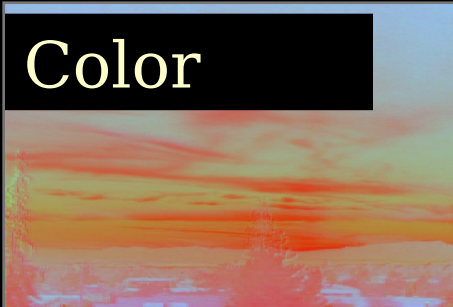
Intensity



Gamma on
intensity



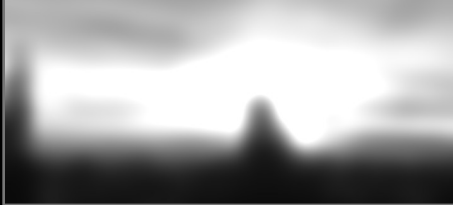
Color



Chiu et al. 1993

- Reduce contrast of low-frequencies
- Keep high frequencies

Low-freq.



High-freq.



Color



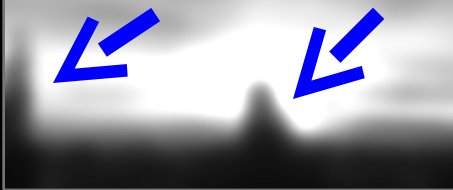
Reduce low frequency



The halo nightmare

- For strong edges
- Because they contain high frequency

Low-freq.



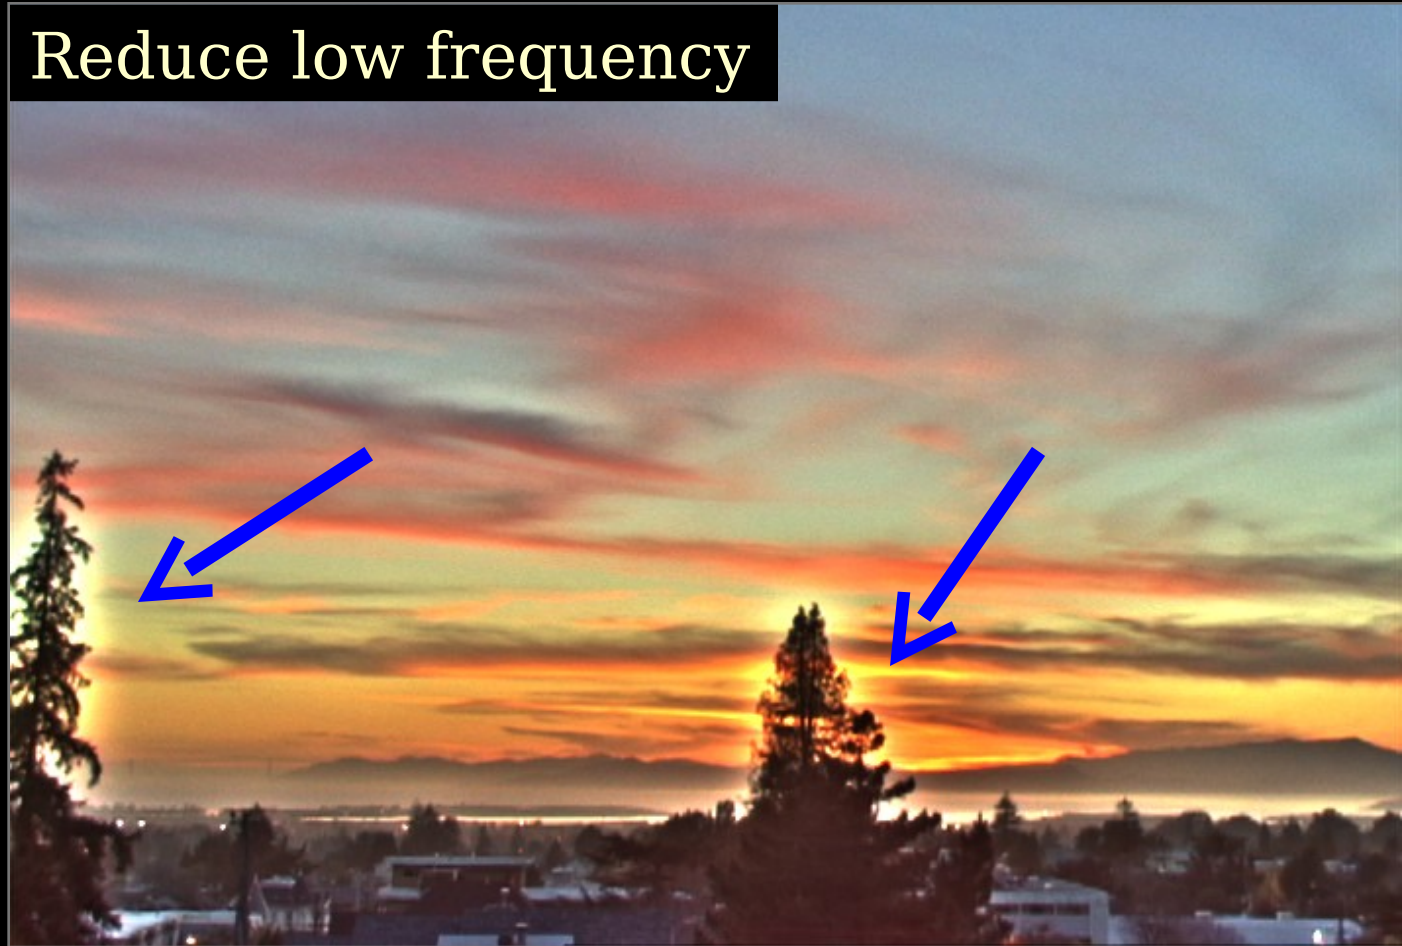
High-freq.



Color



Reduce low frequency



Our approach

- Do not blur across edges
- Non-linear filtering

Large-scale



Detail



Color



Output



Multiscale decomposition

- Multiscale retinex [Jobson et al. 1997]



Compressed

Compressed

Compressed

- Perceptual filters [Pattanaik et al. 1997]



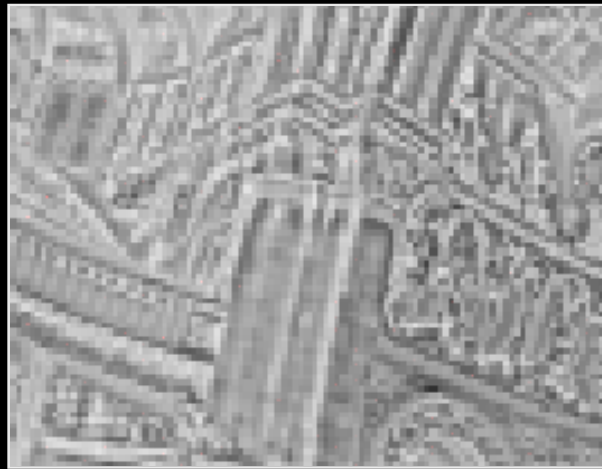
Edge-preserving filtering &

LCIS [Tumblin & Turk 1999]

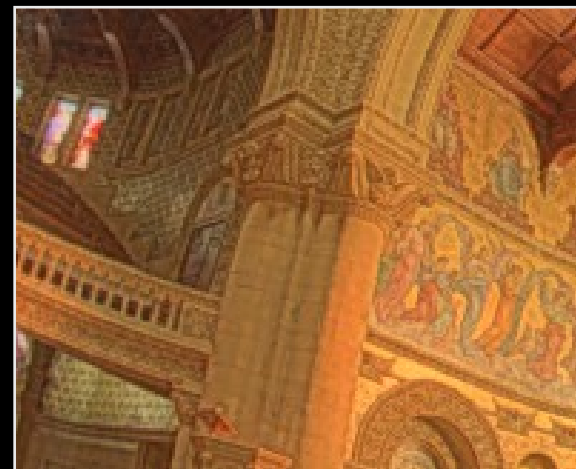
- Multiscale decomposition using LCIS (anisotropic diffusion)



Simplified
(at multiple scales)
Compressed



Detail
s



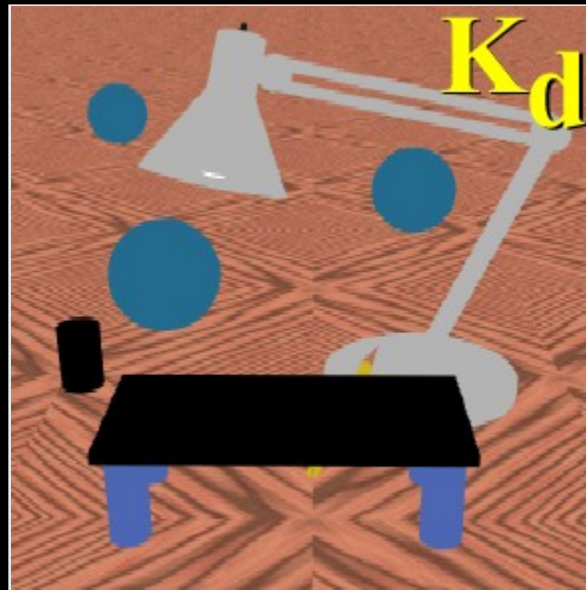
Output

Layer decomposition

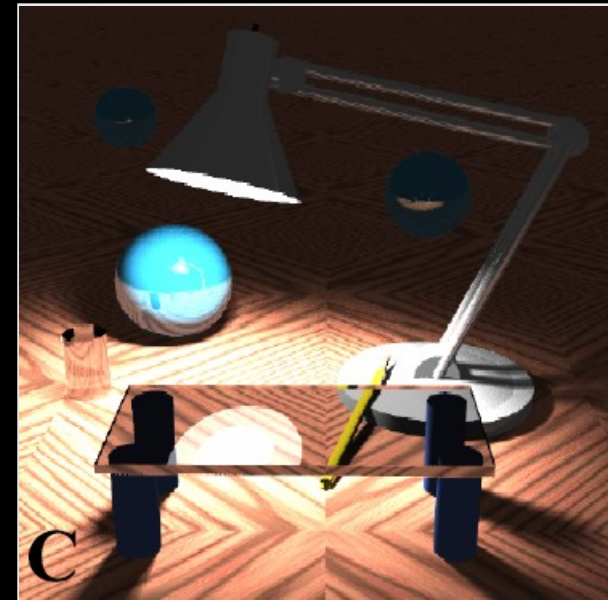
- [Tumblin et al. 1999]
- For 3D scenes
- Reduce only illumination layer



Illumination layer
Compressed



Reflectance layer



Output

Comparison with our

approach

• We use only 2 scales

- Can be seen as illumination and reflectance
- Different edge-preserving filter from LCIS

Large-scale

Detail

Output

Compressed

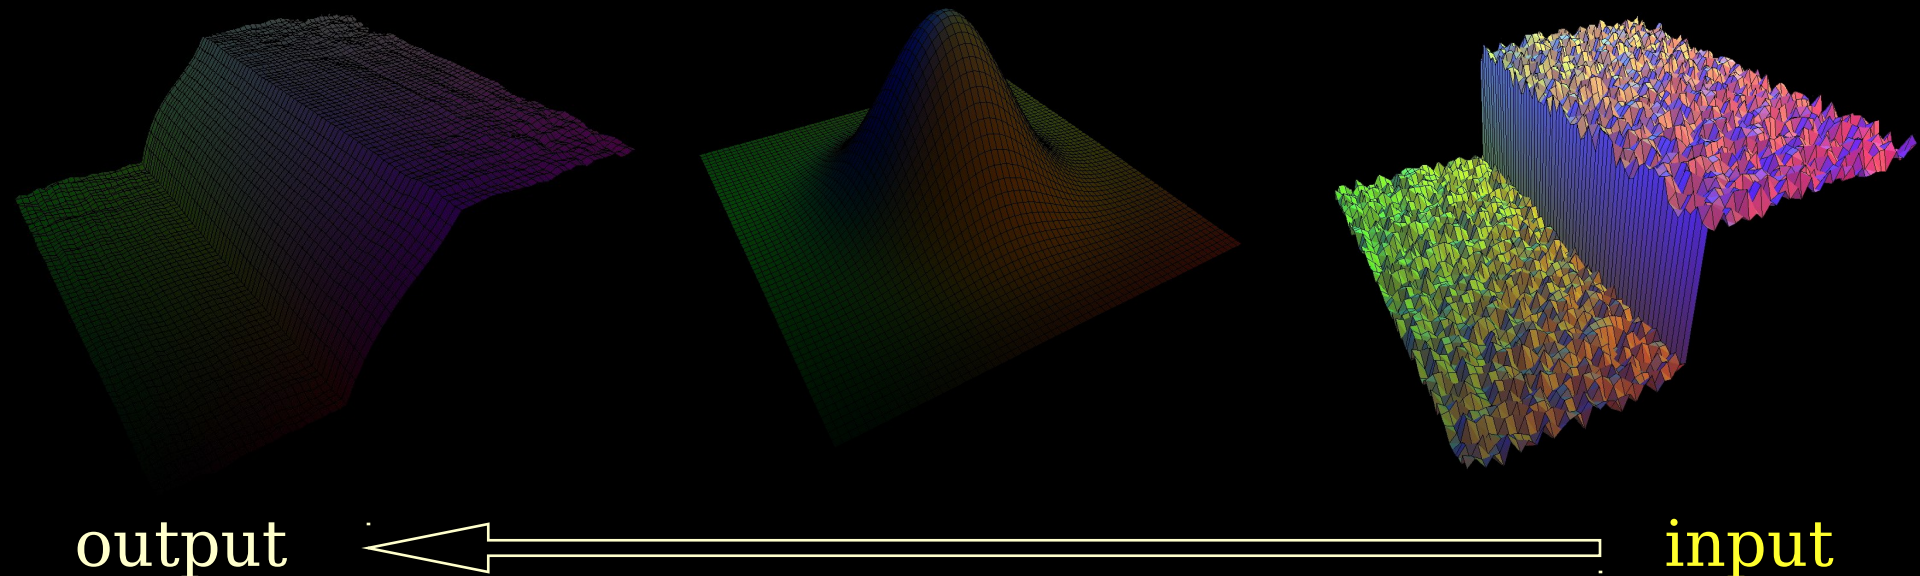


Plan

- Review of bilateral filtering [Tomasi and Manduchi 1998]
- Theoretical framework
- Acceleration
- Handling uncertainty
- Use for contrast reduction

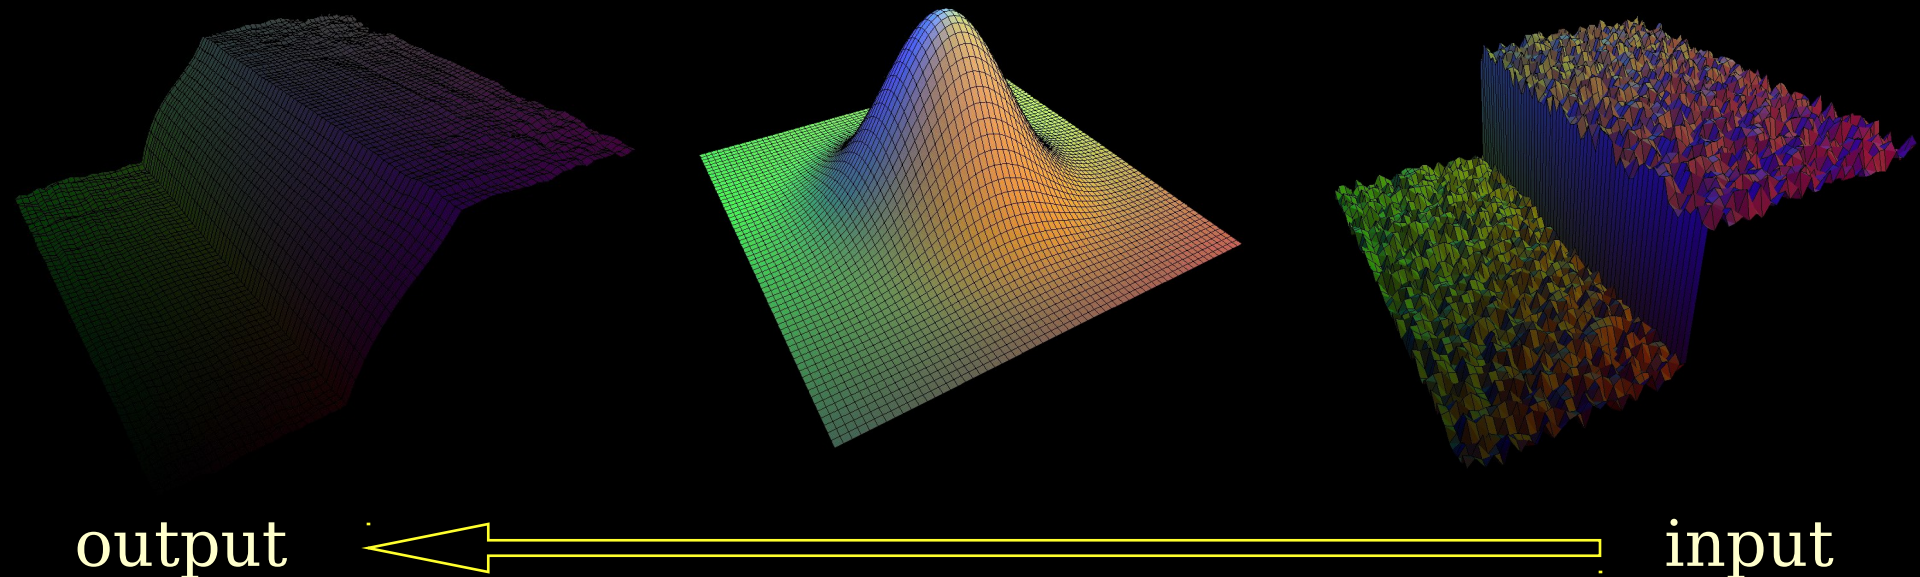
Start with Gaussian filtering

- Here, input is a step function + noise



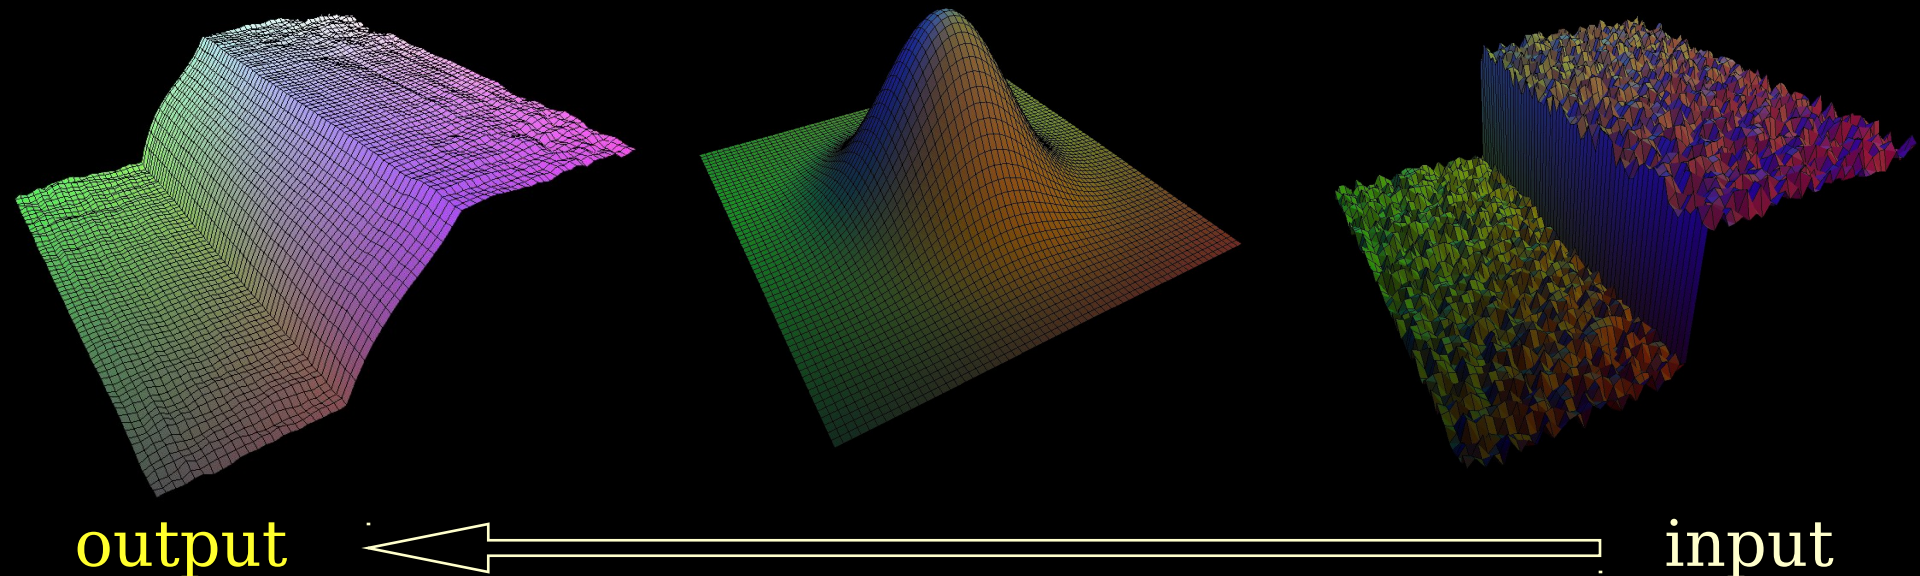
Start with Gaussian filtering

- Spatial Gaussian f



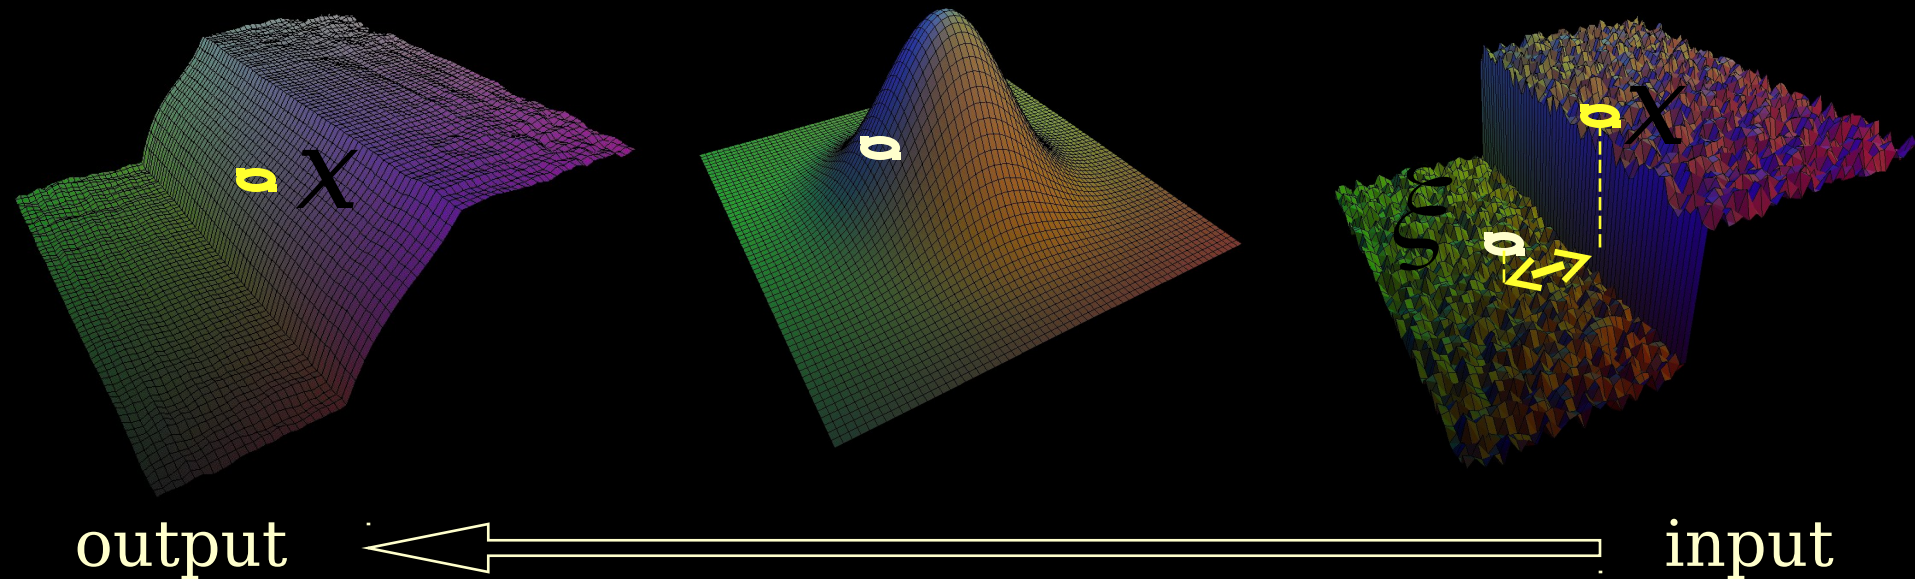
Start with Gaussian filtering

- Output is blurred



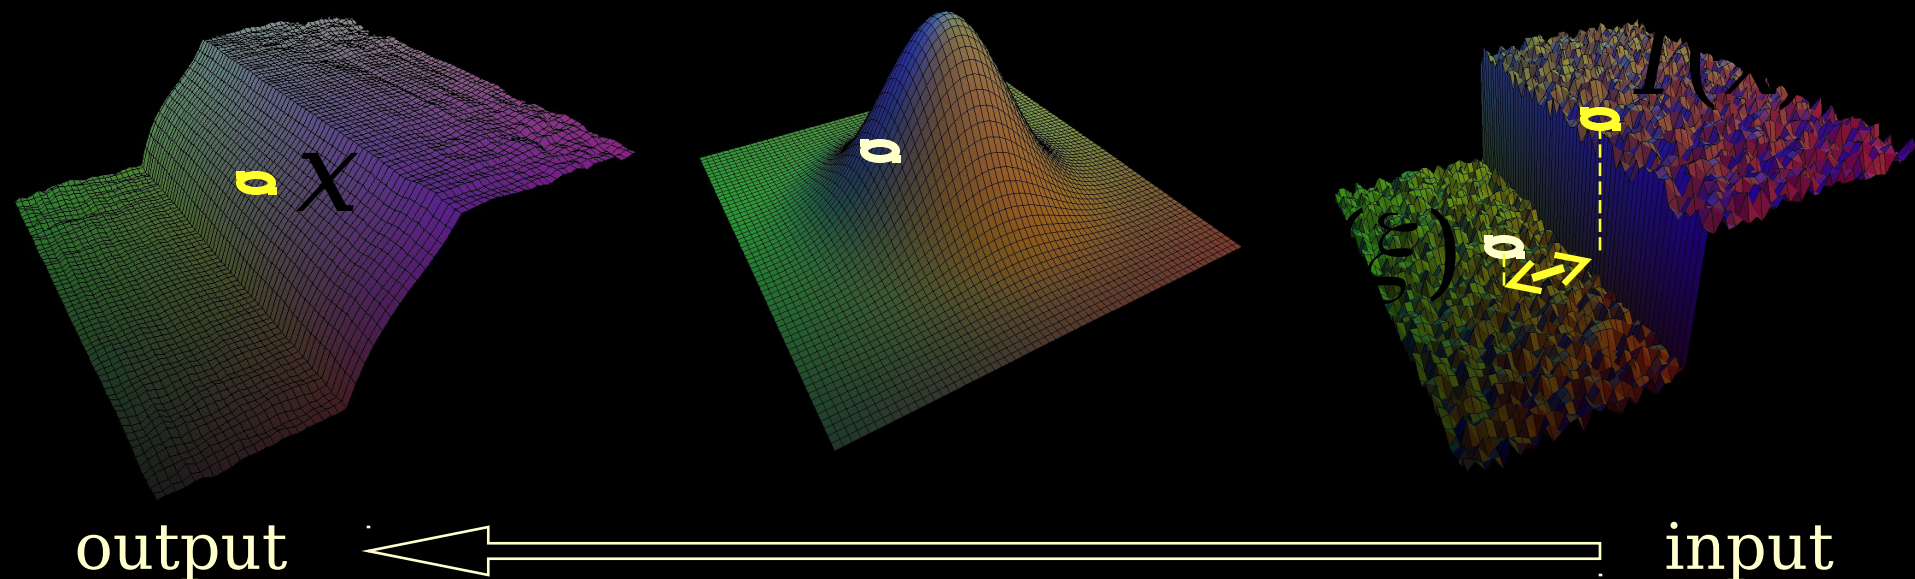
Gaussian filter as weighted

average of ξ depends on distance to x



The problem of edges

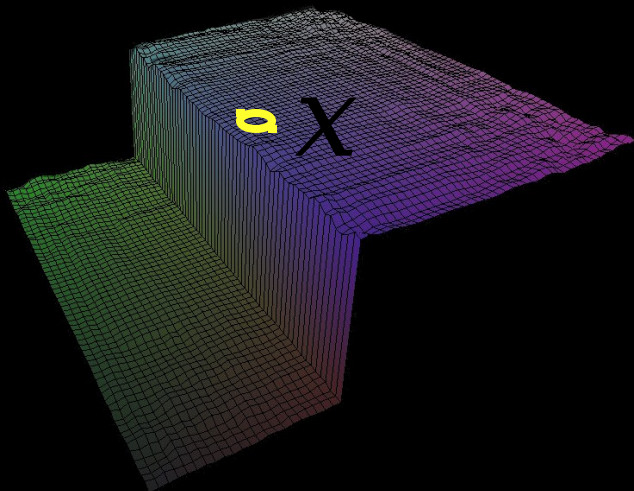
- Here, “pollutes” our estimate $J(x)$
- It is too different



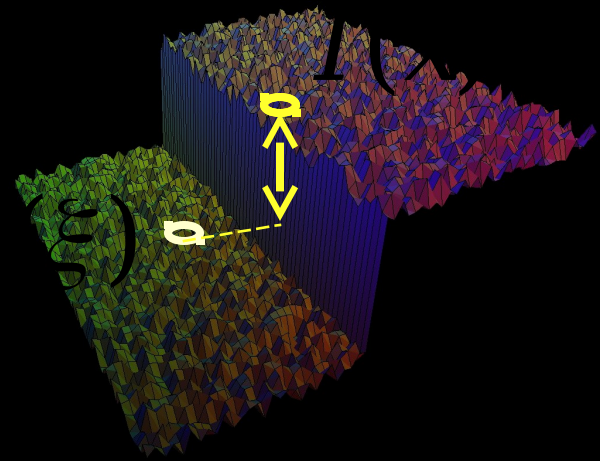
Principle of Bilateral filtering

[Tomasi and Manduchi 1998]

- Penalty g on the intensity difference



output

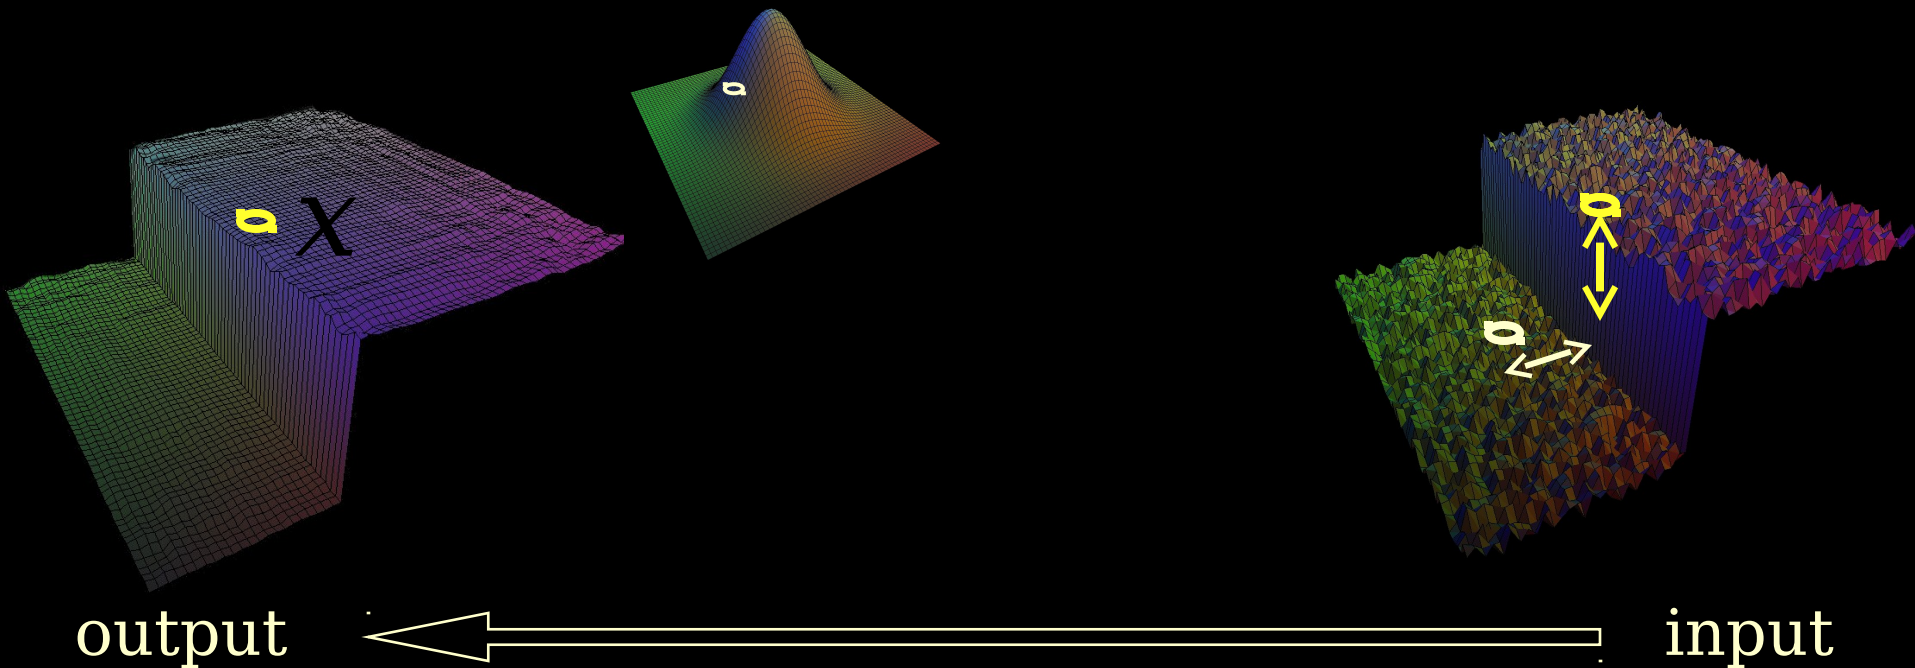


input

Bilateral filtering

[Tomasi and Manduchi 1998]

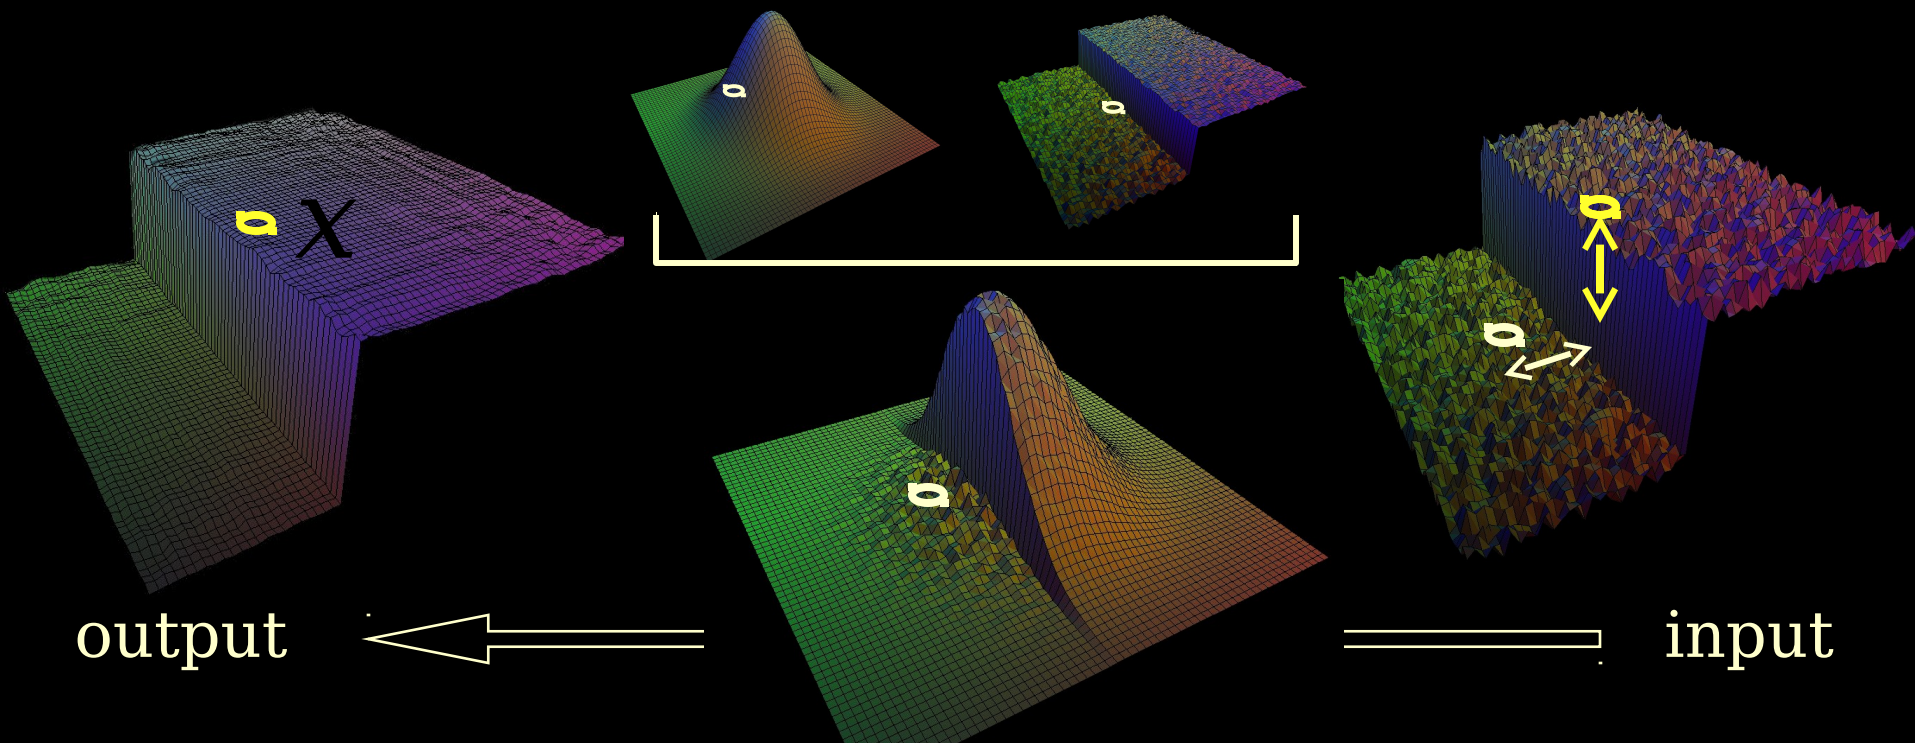
- Spatial Gaussian f



Bilateral filtering

[Tomasi and Manduchi 1998]

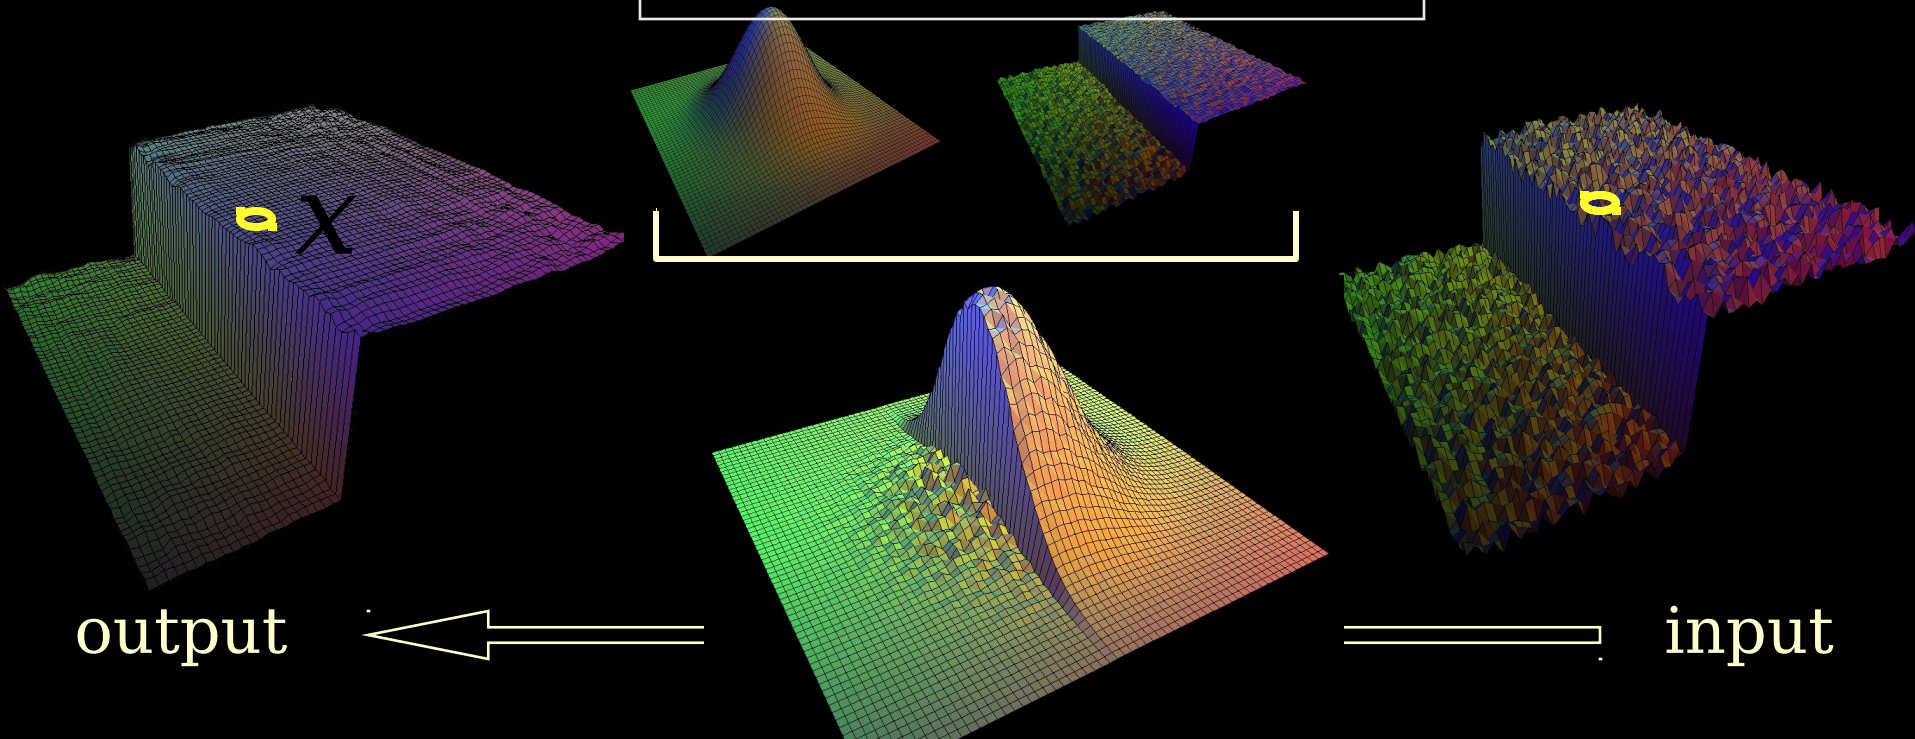
- Spatial Gaussian f
- Gaussian g on the intensity difference



Normalization factor

[Tomasi and Manduchi 1998]

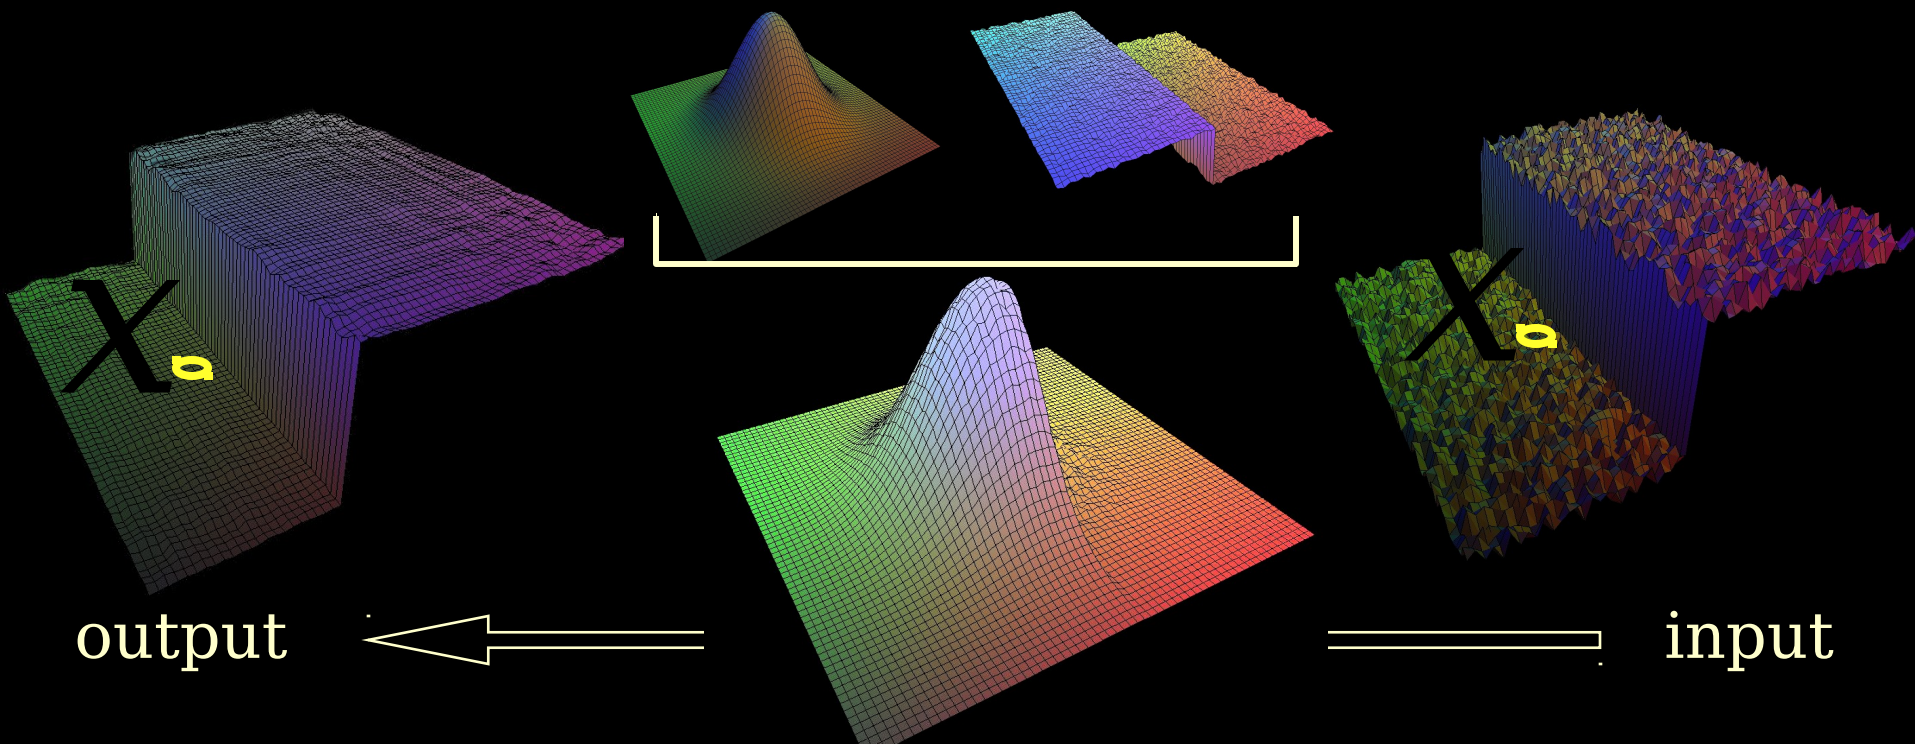
• $k(x) =$



Bilateral filtering is non-

linear [Tomasi and Manduchi 1998]

- The weights are different for each output pixel



Plan

- Review of bilateral filtering [Tomasi and Manduchi 1998]
- Theoretical framework
- Acceleration
- Handling uncertainty
- Use for contrast reduction

Theoretical framework

- Framework of robust statistics
 - Output = estimator at each pixel
 - Less influence to outliers (because of g)
- Unification with anisotropic diffusion
 - Mostly equivalent
 - Some differences
- Details and other insights in paper

Spatial support



Spatial support

- Anisotropic diffusion cannot diffuse across edges



support of anisotropic diffusion

Spatial support

- Anisotropic diffusion cannot diffuse across edges
- Bilateral filtering can



=>

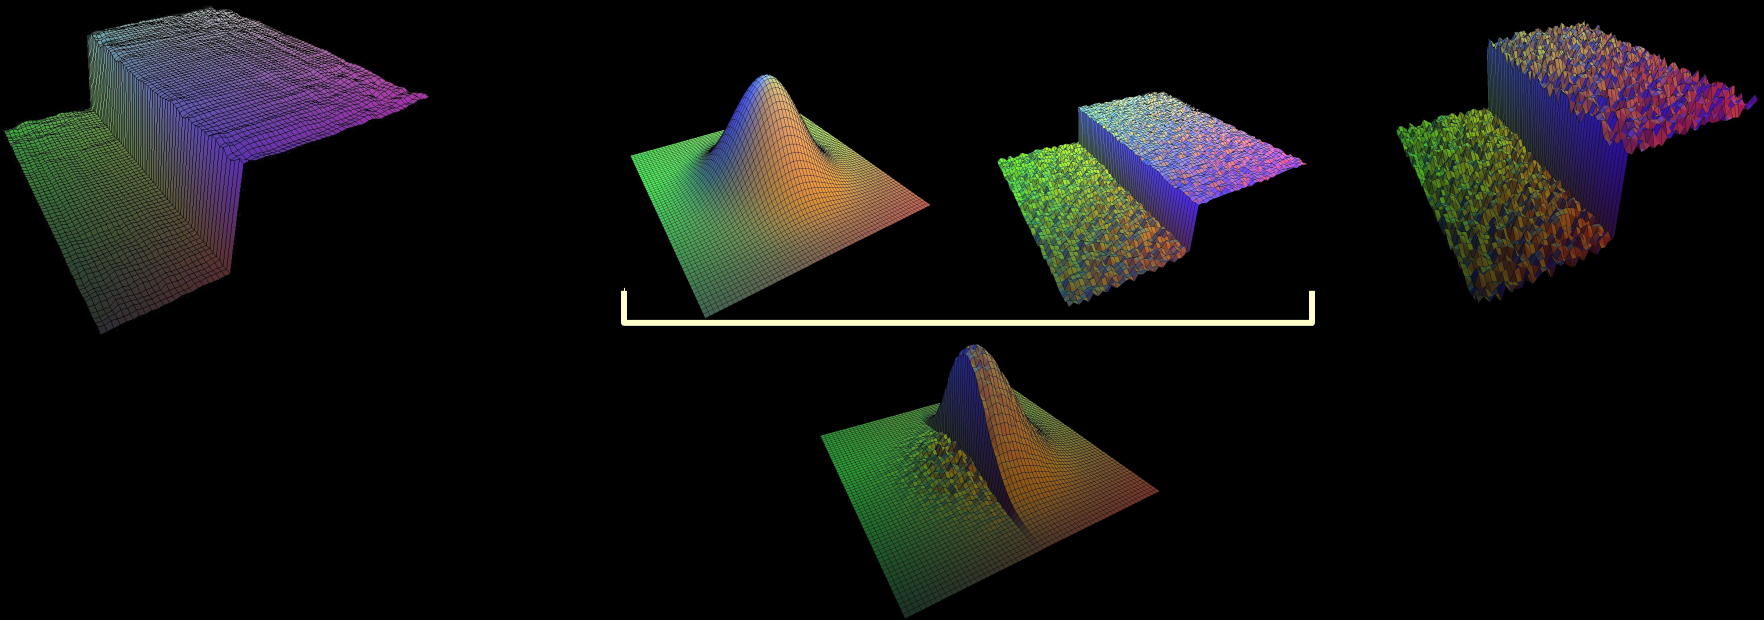


Support of anisotropic diffusion

Support of bilateral

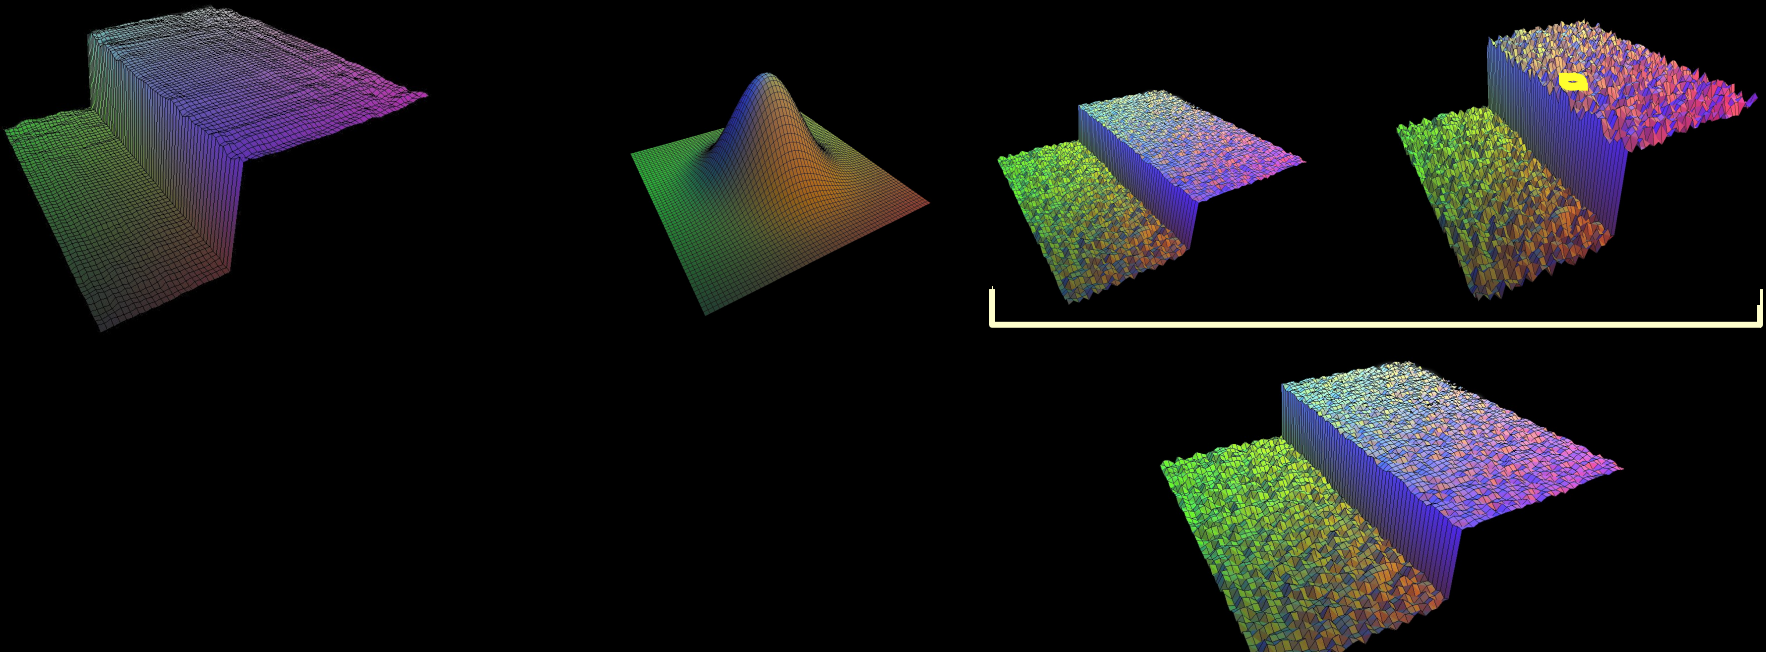
Acceleration

- Non-linear because of g



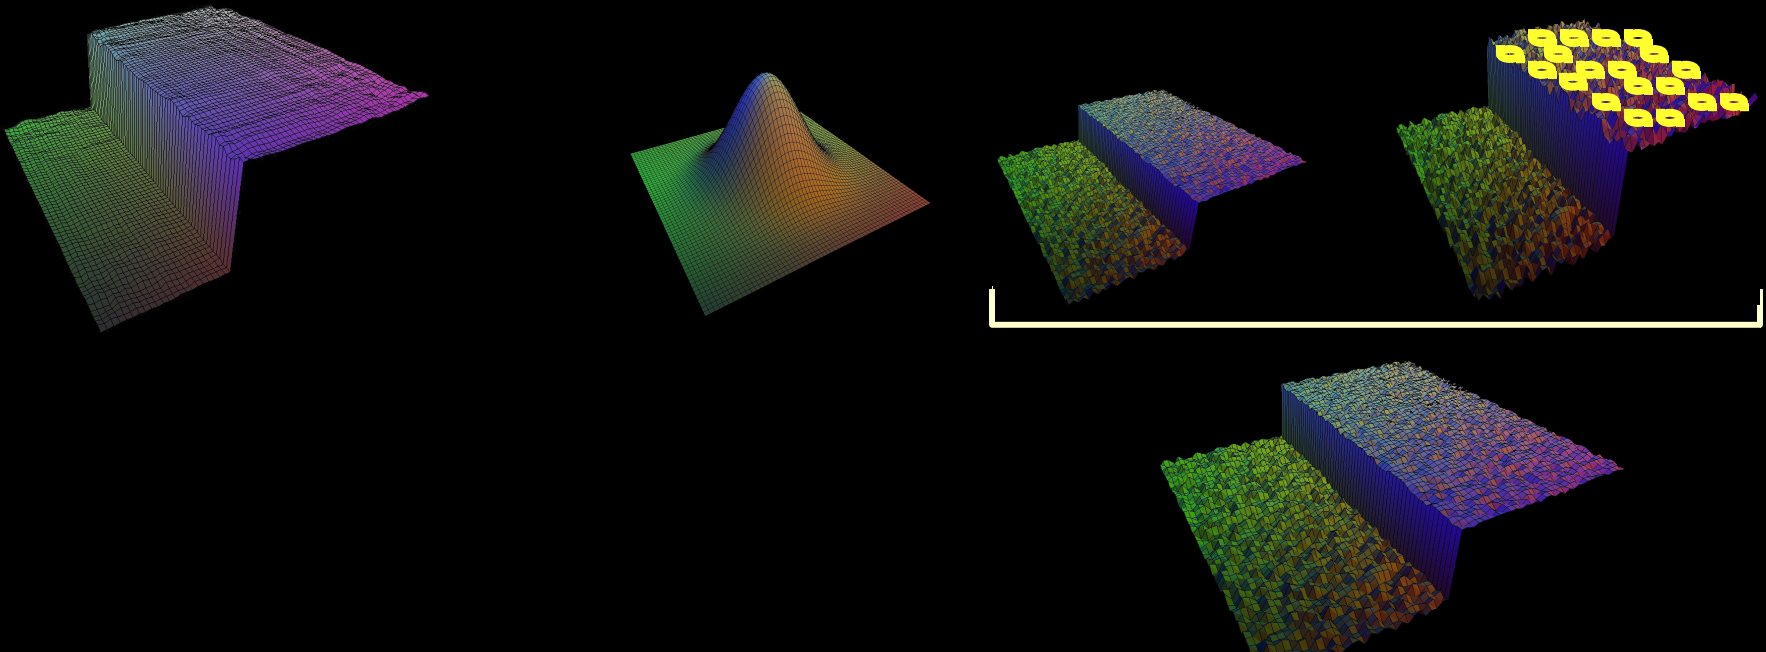
Acceleration

- Linear for a given value of $I(x)$
- Convolution of g I by Gaussian f



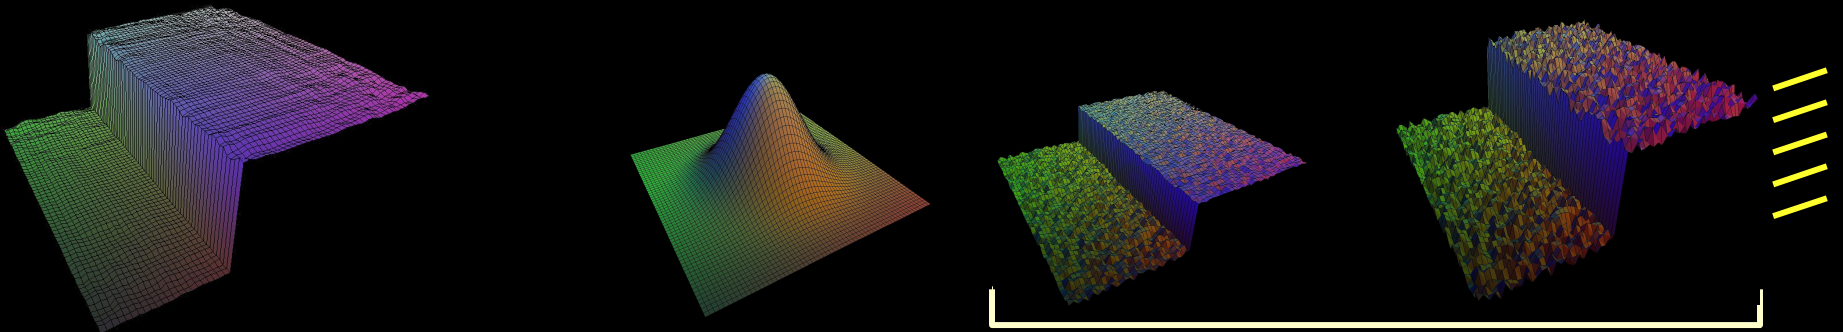
Acceleration

- Linear for a given value of $I(x)$
- Convolution of $g I$ by Gaussian f
- Valid for all x with same value $I(x)$

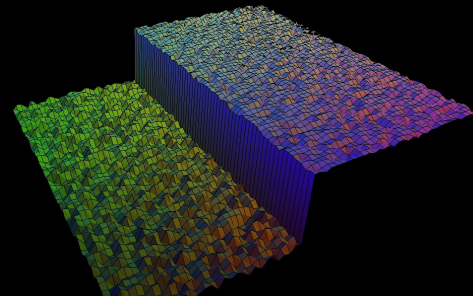


Acceleration

- Discretize the set of possible $I(x)$
- Perform linear Gaussian blur (FFT)
- Linear interpolation in between



- $k(x)$ treated similarly



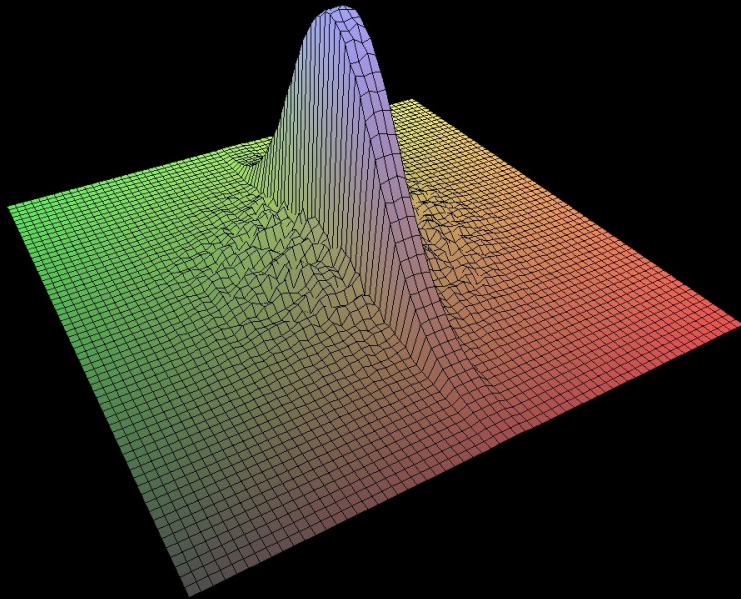
Acceleration

- Piecewise-linearization
 - x10 for a 80pixel kernel on 576*768 image
- Subsampling
 - x30 for a 4x subsampling
 - Superlinear because of cache
- 2 seconds for 2MPixel image (for the complete tone map)



Handling uncertainty

- Sometimes, not enough “similar” pixels
- Happens for specular highlights
- Can be detected using normalization



heights with high uncertainty

ra



Uncertainty

Contrast reduction

Input HDR
image



Contrast
too
high!

Contrast reduction

Input HDR
image



Intensit
y



Color

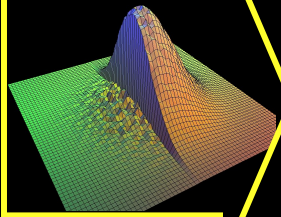


Contrast reduction

Input HDR
image



Intensity



Fast
Bilateral
Filter

Large
scale



Color

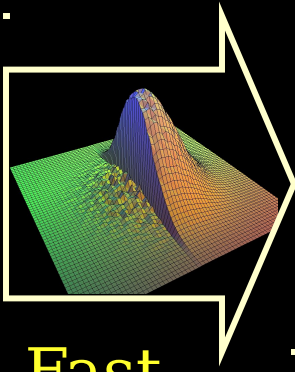


Contrast reduction

Input HDR
image



Intensi
ty



Fast
Bilateral
Filter

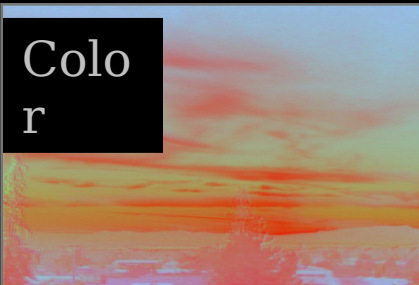
Large
scale



Detail



Colo
r

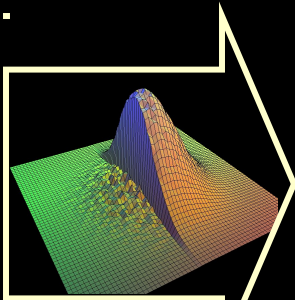


Contrast reduction

Input HDR
image



Intensity



Fast
Bilateral
Filter

Large
scale

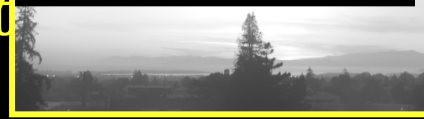


Detail

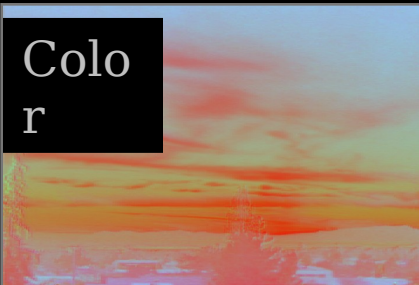


Reduce
contrast

Large
scale

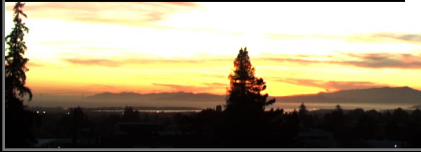


Color

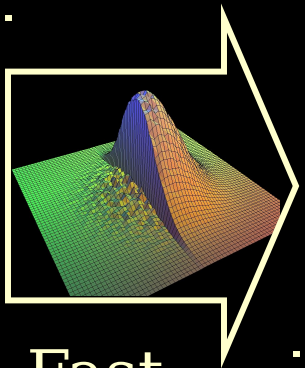


Contrast reduction

Input HDR
image



Intensity



Fast
Bilateral
Filter

Large
scale



Detail



Reduce
contrast
Preserve!

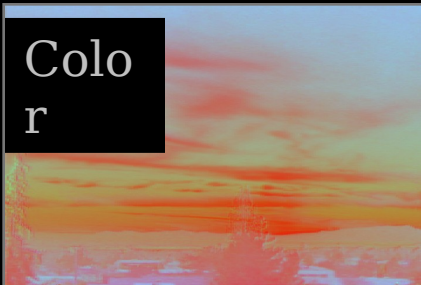
Large
scale



Detail



Color



Contrast reduction

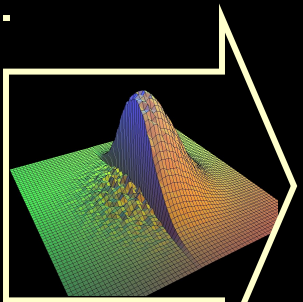
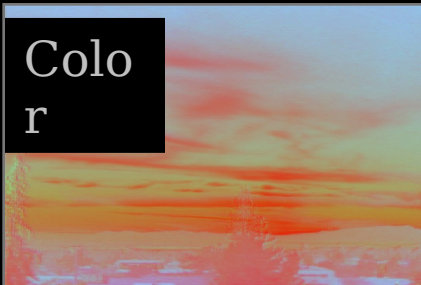
Input HDR image



Intensity



Color



Fast
Bilateral
Filter

Large
scale



Detail



Reduce
contrast
Preserve!

Output



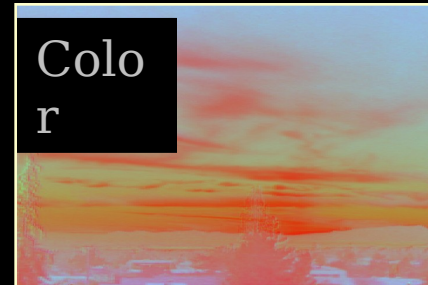
Large
scale



Detail

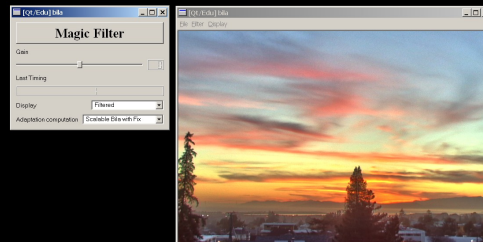


Color



Live demo

- Xx GHz Pentium Whatever PC



Conclusions

- Edge-preserving filter
- Framework of robust statistics
- Acceleration (x300)
- Handling uncertainty

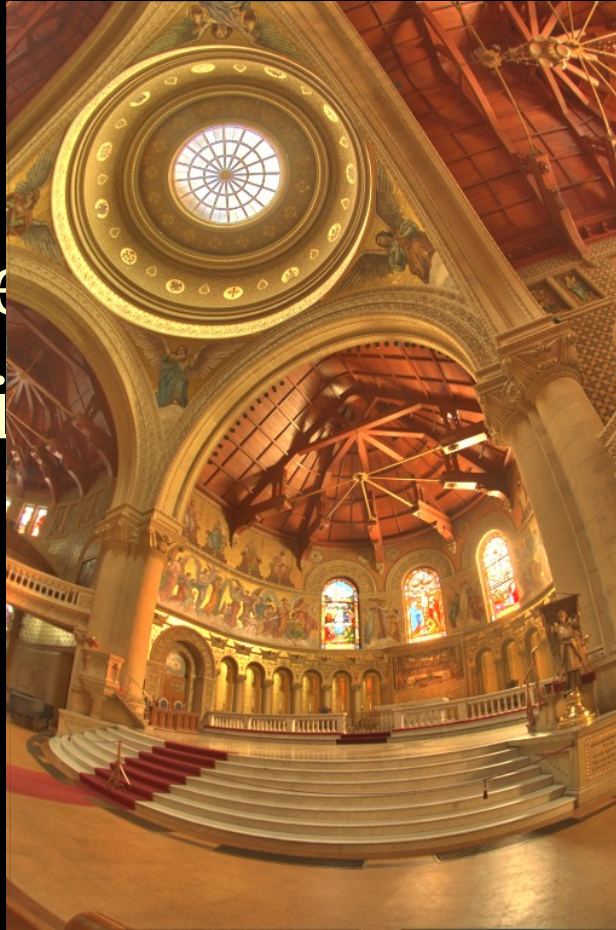
- Contrast reduction
- Can handle challenging photography issues
- Richer sensor + post-processing

Future work

- Uncertainty fix
- Other applications of bilateral filter (meshes, MCRT)
- Video sequences
- High-dynamic-range sensors
- Other pictorial techniques

Acknowledgments

- Mok Oh
- Ray Jones
- Paul Debever
- Jack Tumbli
- Reviewers
- NSF
- Pixar



Informal comparison



Gradient domain
[Fattal et al.]



Bilateral
[Durand et al.]



Photographic
[Reinhard et al.]

Informal comparison



Gradient domain
[Fattal et al.]



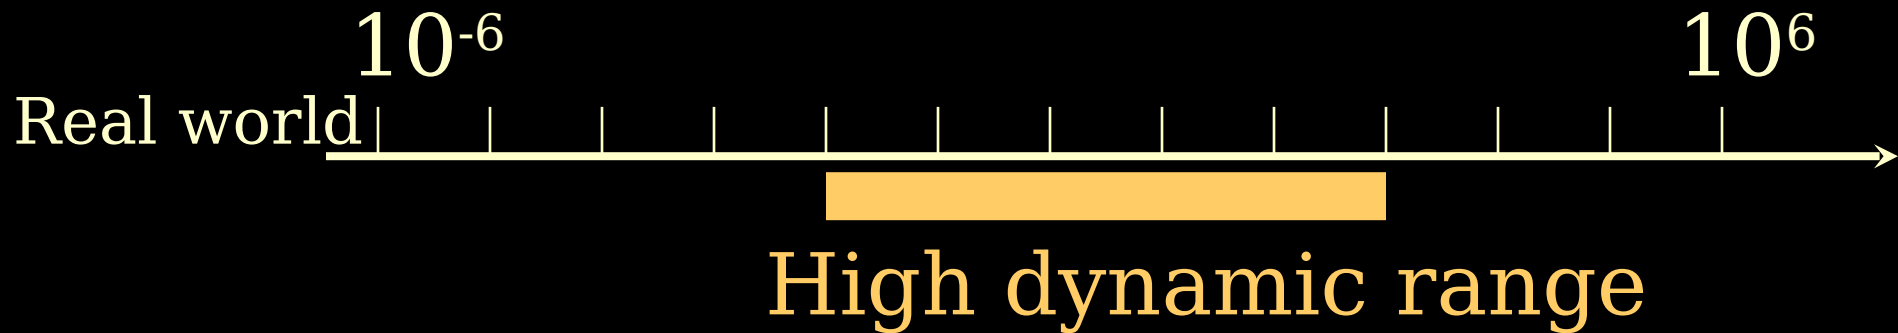
Bilateral
[Durand et al.]





Photographic
[Reinhard et al.]

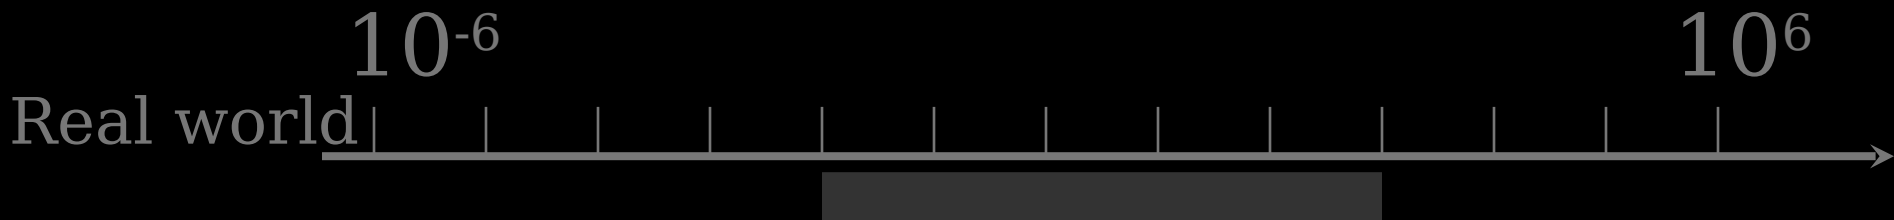
Real world dynamic range

- $\sim 10^{-6}$ to 10^6 cd/m²
- Often 1 : 100,000 in a scene



Picture dynamic range

- Typically 1:50
 - Black  is ~ 50x darker than  white
- Max 1:500



High-dynamic-range (HDR)

images



- Multiple exposure photo [Debevec & Malik 1997]



Recover
response
curve

HDR value
for each
pixel

- HDR sensors



Edge-preserving filtering

- Blur, but not across edges



- Anisotropic diffusion [Perona & Malik 90]
 - Blurring as heat flow
 - LCIS [Tumblin & Turk]
- **Bilateral filtering [Tomasi & Manduci, 98]**

Informal comparison



Gradient-space
[Fattal et al.]



Bilateral
[Durand et al.]



Photographic
[Reinhard et al.]

Informal comparison



Gradient-space
[Fattal et al.]



Bilateral
[Durand et al.]



Photographic
[Reinhard et al.]